



US005548623A

United States Patent [19]

Casper et al.

[11] **Patent Number:** 5,548,623[45] **Date of Patent:** Aug. 20, 1996[54] **NULL WORDS FOR PACING SERIAL LINKS TO DRIVER AND RECEIVER SPEEDS**

5,253,270 10/1993 Petit 375/301

OTHER PUBLICATIONS

[75] Inventors: **Daniel F. Casper**, Poughkeepsie;
Thomas A. Gregg, Highland; **Gregory Salyer**, Woodstock; **Douglas W. Westcott**, Rhinebeck, all of N.Y.

IBM Technical Disclosure Bulletin vol. 31 No. 7 Dec. 1988 pp. 487-491.

"Bit Insertion and Deletion Architecture for Computer-Controlled Data Communications Systems".

IBM Technical Disc. Bulletin vol. 33 No. 3A Aug. 1990 pp. 123-124. "Asynchronous Communications Start and Stop Commands Control the Transmitter of Vart".

[73] Assignee: **International Business Machines Corporation**, Armonk, N.Y.

[21] Appl. No.: 71,150

[22] Filed: Jun. 1, 1993

Primary Examiner—Stephen Chin

Assistant Examiner—T. Ghebretinsae

Attorney, Agent, or Firm—Lynn L. Augspurger; Marc A. Ehrlich

Related U.S. Application Data

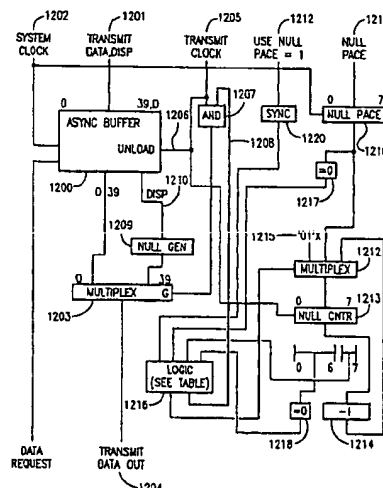

[63] Continuation-in-part of Ser. No. 839,657, Feb. 20, 1992, Pat. No. 5,357,608, and a continuation-in-part of Ser. No. 839,986, Feb. 20, 1992, Pat. No. 5,267,240, and a continuation-in-part of Ser. No. 839,652, Feb. 20, 1992, Pat. No. 5,412,803.

[51] Int. Cl.⁶ **H04L 7/00**[52] U.S. Cl. **375/358; 370/102; 455/69**[58] Field of Search **375/104, 112, 375/358, 377; 370/102, 84; 455/69**[56] **References Cited****U.S. PATENT DOCUMENTS**

3,051,929	8/1962	Smith .	
3,480,734	11/1969	Noriyoshi .	
3,504,287	3/1970	Deregnacourt .	
3,536,840	10/1970	Sullivan	375/58
3,564,414	2/1971	Ebert .	
3,777,066	12/1973	Nicholas	375/118
4,017,688	4/1977	Callens et al. .	
4,025,720	5/1977	Pachynski, Jr. et al. .	
4,151,373	4/1979	Widmer et al. .	
4,330,856	5/1982	Takasaki et al. .	
4,486,739	12/1984	Franaszek et al. .	
4,558,455	12/1985	Edenoy et al. .	
4,661,966	4/1987	Schreiner	375/112
4,979,185	12/1990	Bryans et al. .	
5,003,558	3/1991	Gregg et al. .	
5,048,062	9/1991	Gregg et al. .	

[57] **ABSTRACT**

A system for the transmission of information between elements of a data processing complex and a method for establishing such a system. Two elements of a data processing system are connected by a physical link comprising multiple conductors attached to transceivers at channels in each data processing element. Once the transceivers have been synchronized, commands and responses are exchanged which ensure that all of the transceivers in a channel are connected to the same channel on the other end of the conductor. If the transceivers are considered configured an entry is made in a Configured-Transceiver table. A search is made of an Allowed-Operational-Link table which contains sets of transceivers which are allowed to become operational links. The set of transceivers thus found, is compared against the Configured-Transceiver-Table to verify that all of the members of the set have been configured. If a match is found, this set of transceivers becomes an Intended-Operational-Link. The Intended-Operational-Link is verified to ensure that both channels agree on the set of conductors will form the operational link. If the Intended-Operational-Link verifies, the operational link is established therefrom. Pacing information is exchanged on the link causing null words to be inserted at an optimum rate to compensate for slow receivers.

7 Claims, 12 Drawing Sheets


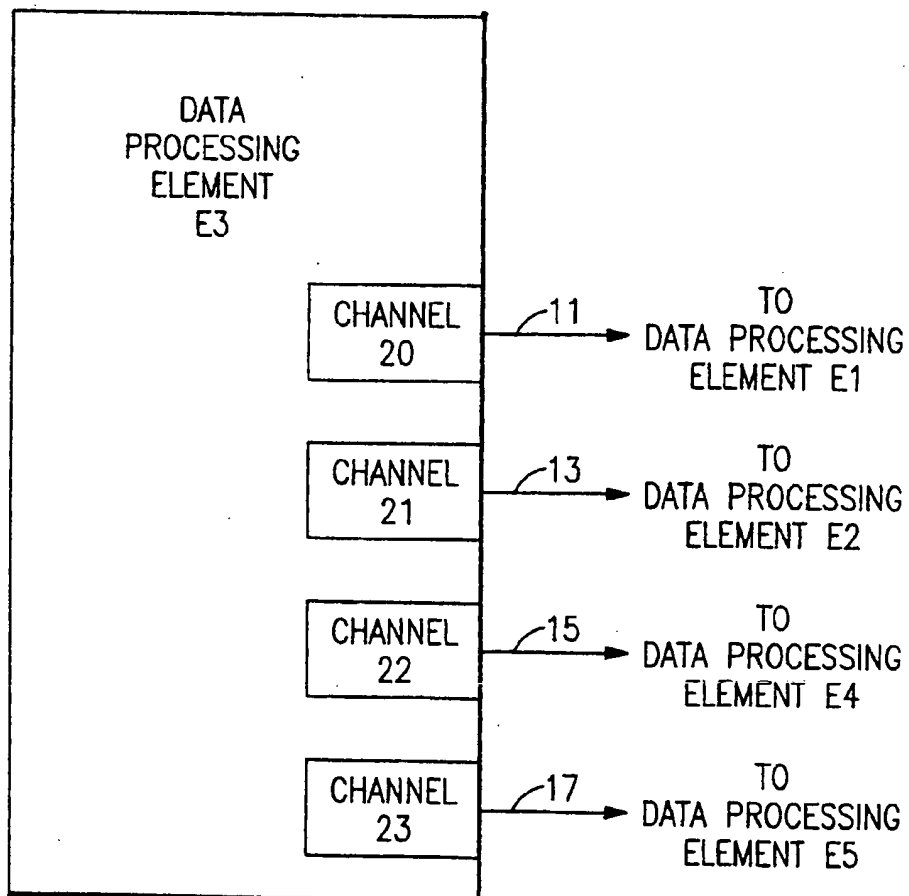

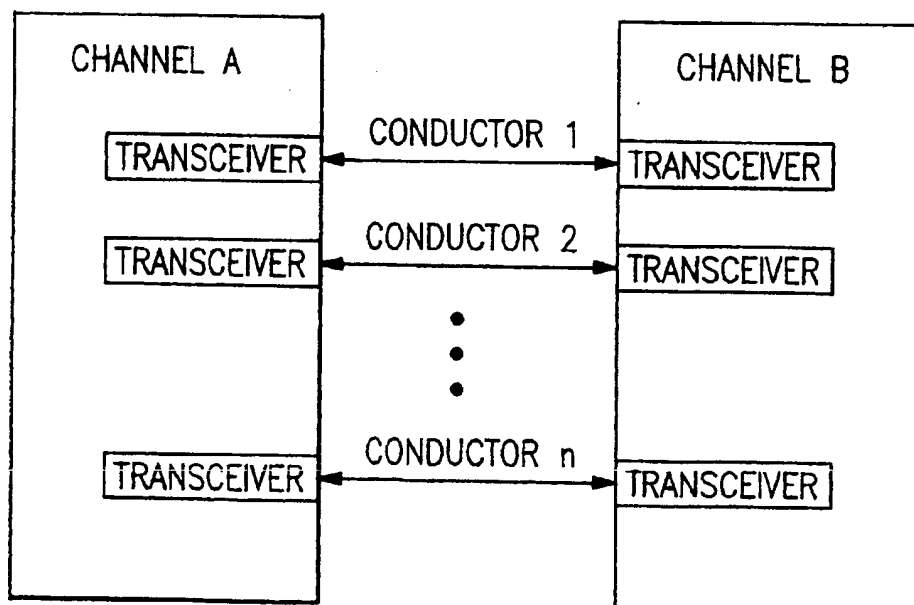


FIG. 1

FIG. 2



FIG. 3FIG. 4

CONFIGURED-TRANSCIVER SUBTABLE

CONDUCTOR	STATE
1	CONFIGURED/NOT CONFIGURED
2	CONFIGURED/NOT CONFIGURED
3	CONFIGURED/NOT CONFIGURED
4	CONFIGURED/NOT CONFIGURED

FIG.5

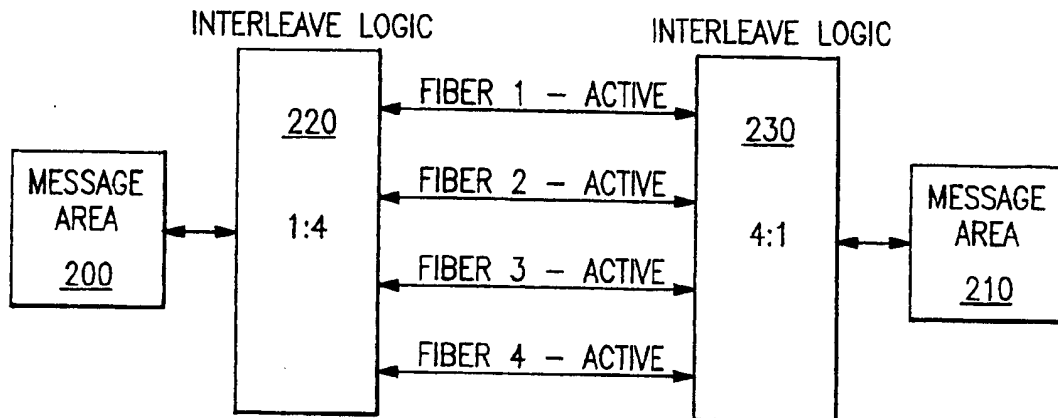
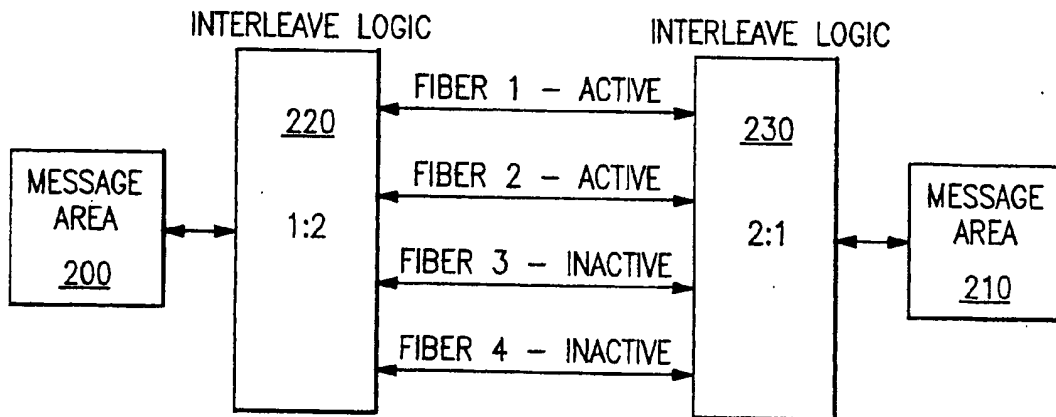
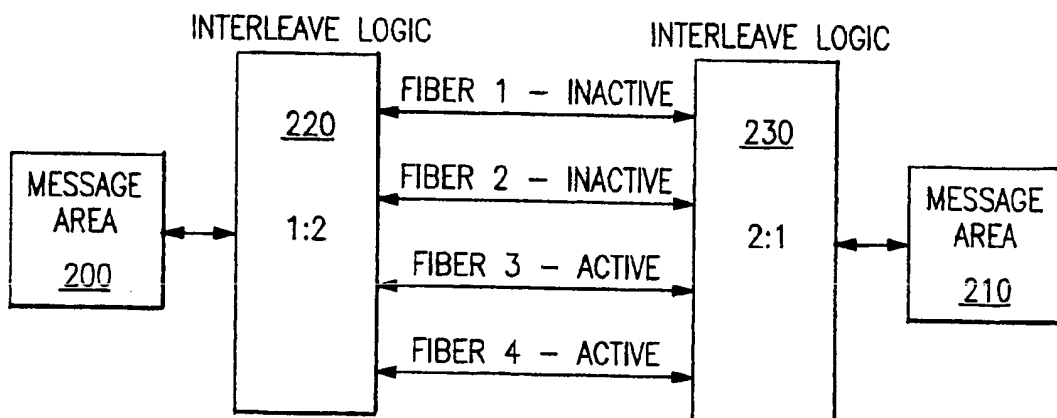
ALLOWED-OPERATIONAL-LINK SUBTABLE

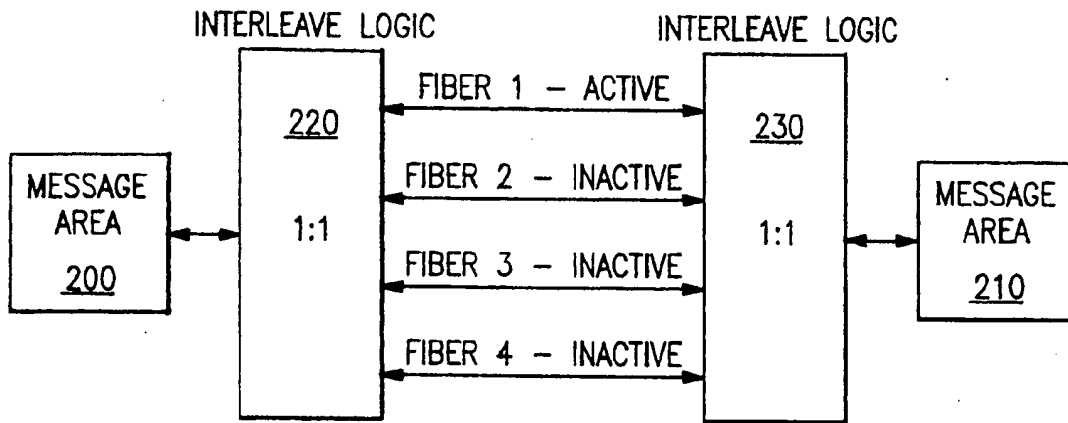
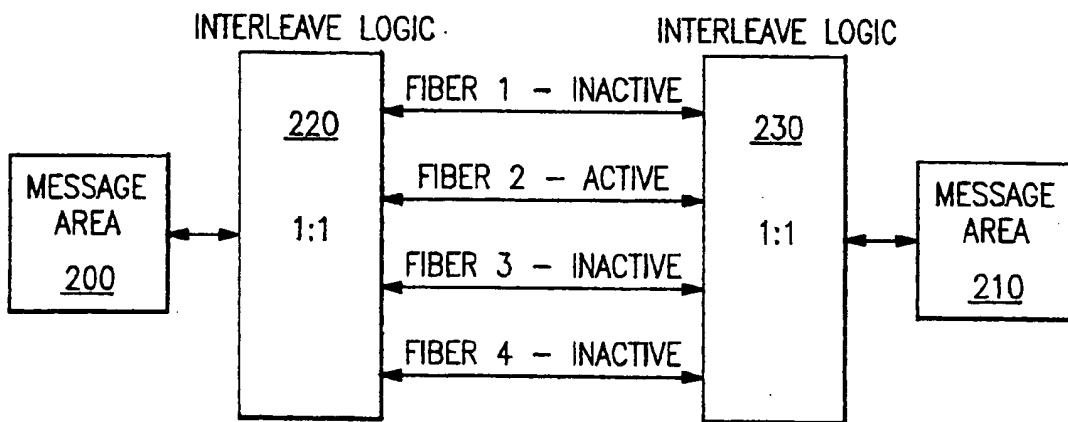
SEARCH ORDER	CONDUCTORS
1	1-2-3-4
2	1-2
3	3-4
4	1
5	2
6	3
7	4

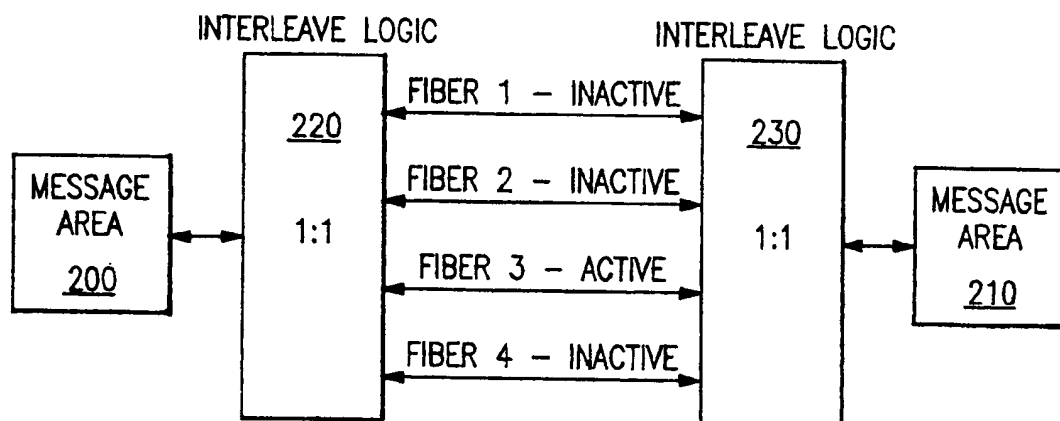
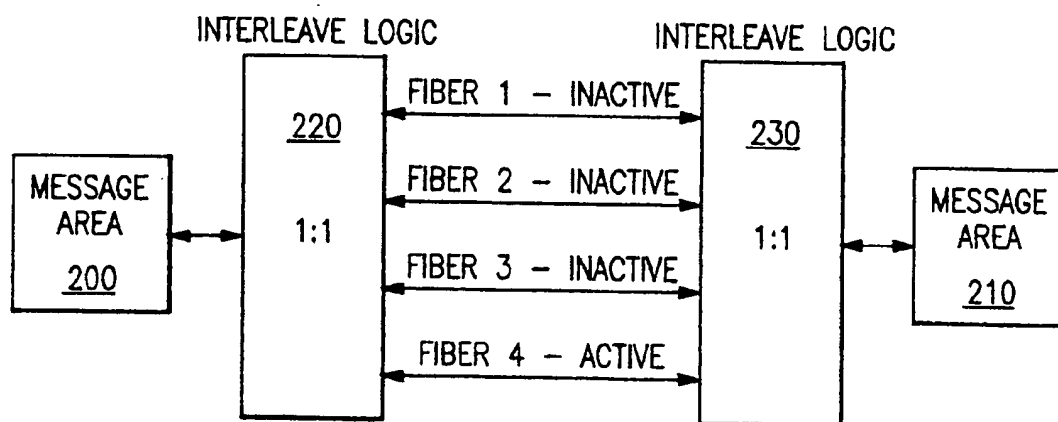
FIG.7

- ALL OF THE TRANSCEIVERS
- THE HIGH ORDER ONE HALF OF THE TRANSCEIVERS
- THE LOW ORDER ONE HALF OF THE TRANSCEIVERS
- THE FIRST QUARTER OF THE TRANSCEIVERS
- THE SECOND QUARTER OF THE TRANSCEIVERS
- THE THIRD QUARTER OF THE TRANSCEIVERS
- THE FOURTH QUARTER OF THE TRANSCEIVERS
- THE FIRST EIGHTH OF THE TRANSCEIVERS
- THE SECOND EIGHTH OF THE TRANSCEIVERS
- THE THIRD EIGHTH OF THE TRANSCEIVERS
- THE FOURTH EIGHTH OF THE TRANSCEIVERS
- THE FIFTH EIGHTH OF THE TRANSCEIVERS
- THE SIXTH EIGHTH OF THE TRANSCEIVERS
- THE SEVENTH EIGHTH OF THE TRANSCEIVERS
- THE EIGHTH EIGHTH OF THE TRANSCEIVERS
- THE FIRST SIXTEENTH OF THE TRANSCEIVERS
- THE SECOND SIXTEENTH OF THE TRANSCEIVERS
-
-
-
- THE FIFTEENTH SIXTEENTH OF THE TRANSCEIVERS
- THE SIXTEENTH SIXTEENTH OF THE TRANSCEIVERS
- THE FIRST THIRTY-SECOND OF THE TRANSCEIVERS
- THE SECOND THIRTY-SECOND OF THE TRANSCEIVERS
-
-
-
- THE THIRTY FIRST THIRTY-SECOND OF THE TRANSCEIVERS
- THE THIRTY SECOND THIRTY-SECOND OF THE TRANSCEIVERS
- THE FIRST SIXTY-FOURTH OF THE TRANSCEIVERS
- THE SECOND SIXTY-FOURTH OF THE TRANSCEIVERS
-
-
-
- THE SIXTY THIRD SIXTY-FOURTH OF THE TRANSCEIVERS
- THE SIXTY FOURTH SIXTY-FOURTH OF THE TRANSCEIVERS

FIG.6

FIG. 8aFIG. 8b(1)FIG. 8b(2)

FIG. 8c(1)FIG. 8c(2)

FIG. 8c(3)FIG. 8c(4)







FIG. 9

FIG.10

**FIG.11**

**FIG. 12**

INPUTS					OUTPUTS			
USE PACE=1	PACE =0	CNTR 0:6=0	CNTR BIT7		UNLOAD BUFFER	DECREMENT NULL CNTR	NULL CNTR = '01'X	NULL CNTR = NULL PACE
0	0	0	0		1	1	0	0
0	0	0	1		1	1	0	0
0	0	1	0		0	0	0	1
0	0	1	1		1	1	0	0
0	1	0	0		1	X	X	X
0	1	0	1		1	X	X	X
0	1	1	0		1	X	X	X
0	1	1	1		1	X	X	X
1	0	0	0		1	0	1	0
1	0	0	1		1	0	1	0
1	0	1	0		0	0	1	0
1	0	1	1		1	1	0	0
1	1	0	0		1	0	1	0
1	1	0	1		1	0	1	0
1	1	1	0		1	0	1	0
1	1	1	1		1	1	0	0

FIG.13

NULL WORDS FOR PACING SERIAL LINKS TO DRIVER AND RECEIVER SPEEDS

This application is a continuation-in-part of U.S. patent application Ser. No. 07/839,657, now U.S. Pat. No. 5,357, 608 filed Feb. 20, 1992 of Neil G. Barrow et al entitled "Configurable, Recoverable Parallel Bus" (IBM Docket No. PO9-91-066).

This application is also related to and is a continuation-in-part of U.S. patent application Ser. No. 07/839,986, now U.S. Pat. No. 5,267,240 filed Feb. 20, 1992 entitled "Frame Group Transmission and Reception for Parallel/Serial Buses" (IBM Docket No. PO9-92-001).

This application is also related to and is a continuation-in-part of U.S. patent application Ser. No. 07/839,652, now U.S. Pat. No. 5,412,803 filed Feb. 20, 1992 entitled "High Performance Channels for Data Processing Systems" (IBM Docket No. PO9-91-067).

FIELD OF THE INVENTION

This invention generally relates to the field of data processing, more specifically to an apparatus and method of communication between elements of a data processing complex to allow accommodation of an element which operates at a speed higher or lower than that of the link.

CROSS REFERENCE TO RELATED APPLICATIONS

This application is related to the following applications filed concurrently herewith:

The continuation-in-part of U.S. patent application Ser. No. 07/839,657, filed Feb. 20, 1992 of Neil G. Barrow et al entitled "Configurable, Recoverable Parallel Bus" (IBM Docket No. PO9-91-066X), now U.S. Ser. No. 08/071,146, filed Jun. 1, 1993.

The continuation-in-part of U.S. patent application Ser. No. 07/839,986 filed Feb. 20, 1992 entitled "Frame Group Transmission and Reception for Parallel/Serial Buses" (IBM Docket No. PO9-92-001X), now U.S. Ser. No. 08/071,115, filed Jun. 1, 1993.

U.S. Ser. No. 08/071,154, filed Jun. 1, 1993, by Neil G. Barrow et al entitled "Quiesce and Unquiesce Function for Intersystem Channels" (IBM Docket No. PO9-93-015);

U.S. Ser. No. 08/070,587, filed Jun. 1, 1993, by Kenneth J. Fredericks et al entitled "Concurrent Maintenance of Degraded Parallel/Serial Buses" (IBM Docket No. PO9-93-016);

U.S. Ser. No. 08/070,589, filed Jun. 1, 1993, by Thomas A. Gregg et al entitled "Error Detection and Recovery in Parallel/Serial Buses" (IBM Docket No. PO9-93-018).

These co-pending applications and the present application are owned by one and the same assignee, International Business Machines Corporation of Armonk, N.Y.

The descriptions set forth in these co-pending applications are hereby incorporated into the present application by this reference.

BACKGROUND OF THE INVENTION

Fiber optic waveguides enable the serial transmission of information from a driver to a receiver at a relatively long distance (kilometers) and at very high data rates (billions of bits per second). High performance data processing complexes require system buses among its elements to have even higher bandwidths of information transfer (100's of millions

of bytes per second). The use of a single conductor, such as a fiber optic waveguide, can be insufficient to serve as such a system bus because of bandwidth limitations.

In the prior art, one solution to the bandwidth problem has been to use multiple conductors in parallel, each conductor carrying a bit of each word to be transmitted. A second solution, known as data striping, entails sending a plurality of data words in parallel over multiple conductors in a serial fashion. Each word is transmitted serially over a single conductor, but it is transmitted in parallel with other words on other conductors.

One challenge in implementing a data striping technique for information transmission is the determination of which of the multiple conductors attached to a computer element form the system bus linking one computer element to another. Some prior art systems attempt to use each of the conductors which is physically attached to the computer element. This method has a significant drawback in that if one of the conductors is not operational, the entire bus is rendered inoperable. Other prior art systems have attempted to configure a system bus using only the operational conductors, but these prior art systems have been driven by hardware switches and have been further limited to system buses of two conductors.

When a single conductor has the ability to operate at speeds greater than the device driving the connection, or receiving the signal, a method of pacing the serial connection to the devices at each end of the connection is required. This allows the maximum bandwidth to be derived from the use of each conductor. Prior art acids pacing bits only based on the senders system clock rate difference with respect to the conductor's clock rate. Other prior art requires the messages to be fully buffered at the sender prior to the transmission of the message over the conductor at the conductor's clock rate.

SUMMARY OF THE INVENTION

The present invention is a system for the transmission of information between elements of a data processing complex and a method for establishing such a system. Each element of the data processing complex contains at least one channel for enabling communication to other elements of the complex. The channel incorporates the controls, memory and physical connections required for the communication. A physical link, comprising multiple conductors, is the means by which any two channels are connected. In the preferred embodiment, these conductors are one or more pairs of optical fibers. The fiber-pairs are connected to a channel by transceivers. The transceivers enable data to flow in one direction on one of the fibers in the pair, and in the opposite direction on the other fiber. The method of the present invention determines which of the multiple pairs of fibers actually form an operational link between the two data processing elements. An operational link is the set of transceivers and fibers across which the information that comprises a high level operation is communicated. Because the physical elements of the link, the transceivers and/or fibers, can break or otherwise become inoperable, the method of the present invention also allows the operational link of multiple fiber pairs to be reconfigured at any time, with fewer than a full complement of fibers. The first step in the method is to synchronize the transceivers. When the transceivers have been synchronized, or after the expiration of a specific period of time, a command is issued on each transceiver to acquire a node descriptor for the channel attached to the

other side of the fiber. The node descriptor will uniquely identify the channel and transceiver position to which the transceiver is connected (e.g. this transceiver is connected to central processing complex 1, channel 3, transceiver position 2). After the transceivers have acquired the node descriptors for the other end of the link, the acquired node descriptors are compared to ensure that all of the fibers are connected to the same data processing element. If the node descriptor comparison is valid, then the transceiver is considered configured and all entry is made in a Configured-Transceiver table. As a result of the synchronization and node descriptor process, an entry exists in the Configured-Transceiver table for each transceiver attached to the channel. The entries in this table will indicate whether or not any particular transceiver is configured. If a transceiver has been configured, it can be said to be operationally connected to its counterpart transceiver on the other end of the link. In order to establish the operational link, an Allowed-Operational-Link table is searched. This table contains sets (or combinations) of transceivers which are allowed to become operational links (some combinations of transceivers are not allowed). The table is organized so that the search will first result in a set of transceivers incorporating the greatest number of transceivers connecting the two channels (i.e. the link with the greatest bandwidth). The set of transceivers thus found, is compared against the Configured-Transceiver table to verify whether all of the members of the set have been configured. If even one of the transceivers is not configured, the search of the Allowed-Operational-Link table will continue until a set of fully configured, allowable transceivers has been found. Once a match is found, this set of transceivers becomes an Intended-Operational-Link. As an alternative to using a table, the Allowed-Operational-Links can be dynamically generated by a software program in the system. This software program would successively generate a different Allowed-Operational-Link which would then be tested using the Configured-Transceiver table in order to create the Intended-Operational-Link.

There are three alternative methods by which an Intended-Operational-Link can become the operational link. In a first method, commands are issued from both ends of the link in a peer to peer relationship. In the second method and third method, only one end of the link issues commands, in a master/slave relationship. With respect to the first method, once the Intended-Operational-Link has been determined, each channel issues a command on each transceiver in its Intended-Operational-Link. The channel on the other side of the link which receives this command will compare the received Intended-Operational-Link to its own Intended-Operational-Link in order to verify that there is correspondence. If there is one-to one correspondence, then the recipient channel will send back a response, on each of the transceivers in the Intended-Operational-Link. This response will indicate that the Intended-Operational-Link is acceptable as an operational link. Both ends of the link issue the command on each transceiver in its Intended-Operational-Link, and the responses must be received on the same set of transceivers which sent the command.

The second method for establishing an operational link from an Intended-Operational-Link involves commanding from only one side of the link. With the Intended-Operational-Link determined, one channel, a master, issues a command on each transceiver in its Intended-Operational-Link. The channel which receives this command, a slave, will examine the set of transceivers which received the command to see if they comprise an acceptable operational link. If this set of transceivers comprise a fully configured,

allowed operational link as defined by the Configured-Transceiver-Table and the Allowed-Operational-Link-Table, then the channel which is acting as a slave will send back an appropriate response on each of the transceivers in the acceptable Intended-Operational-Link. The master channel which issued the command examines the responses to ensure that responses were received on all of the transceivers in its Intended-Operational-Link (and no additional transceivers). If there is one-to-one correspondence for the responses, then the Intended-Operational-Link becomes the actual operational link, and link operations can start.

The third method for establishing an operational link for an Intended-Operational-Link involves commanding from only one side of the link. One channel, a master, issues an establish-link command on each transceiver in the transceiver configured state. The frames containing the command have a field which indicates the capability of the transceiver to become a member of the intended-operational-link at the sender of the frame. The channel which receives this command, a slave, will examine the frames and determine the set of transceivers which can become an operational-link. A logical AND operation is performed on a transceiver basis for the state of the transceiver at the channel as determined by the transceiver-configured table entry and the capability of the transceiver to become a member of an intended-operational-link as indicated in the frame received by the slave on the respective transceiver. The results of this logical operation is performed for each of the transceivers and the results used as the candidate transceivers in a search of the allowed-operational-link table and a response is sent from the slave to the master indicating the results on each transceiver in the transceiver-configured state. The response frames returned contain a bit that is unique on each transceiver frame that indicates that the transceiver is part of the intended-operational-link at the slave. Once the slave sends these response frames, the intended-operational link becomes the operational link at the slave. Once the master receives these responses, the intended-operational-link as defined by the received frames is checked against the allowed-operational-link table entries to ensure that the intended-operational-link is valid, and if it is valid, the intended-operational-link becomes the operational-link and use of the link can begin. The master must receive responses on all transceivers in the transceiver configured state before the intended-operational-link can be checked. If responses are not received within a specified time limit, the requests are resent. The allowed-operational-link tables must be the same at both ends of the link for this method to function properly.

In a system where the conductor transmission speed exceeds the capability of the transmitter and/or the receiver, various methods of bit stuffing and deletion have been used. Since the channel described transmits and receives data using the 8/10 code (Franzcek and Widmer) on word (40 bit) boundaries, bit stuffing and deletion is extended to word boundaries.

A method of inserting and deleting bit stuffing words (null words) is provided which does not compromise the transmission and error detection qualities of the 8/10 code, and which can be easily implemented at the interfaces between low speed channel logic and high speed serializer/deserializer logic.

Additionally, a method of exchanging a null pacing parameter between the end points of the conductor instructs the transmitters to insert null words at an optimum rate. The null word insertion rate is computed to be high enough to prevent receiver overrun while not excessively reducing the bandwidth of the conductor.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a generalized depiction of a data processing complex.

FIG. 2 shows a single generalized data processing element and the connections thereto.

FIG. 3 depicts the connections between two channels of the system.

FIG. 4 illustrates a multiconductor link between two channels of the system.

FIG. 5 is a sample Configured-Transceiver table.

FIG. 6 illustrates the organization of the Allowed-Operational-Link table.

FIG. 7 is a sample Allowed-Operational-Link table.

FIG. 8a depicts a four fiber-pair link in which all four fiber-pairs are active.

FIG. 8b(1-2) depict a four fiber-pair link in which only two of the fiber pairs are active.

FIG. 8c(1-4) depict a four fiber-link in which only one of the four fiber-pairs is active.

FIG. 9 illustrates the basic channel buffer structure.

FIG. 10 illustrates a multimessage channel buffer structure.

FIG. 11 illustrates the receive logic required for null word deletion.

FIG. 12 illustrates the transmit logic required for null word insertion.

FIG. 13 illustrates the transmit logic table for null words.

DESCRIPTION OF THE PREFERRED EMBODIMENT

In order for a data processing complex to function effectively, there needs to be an efficient, ideally transparent, system for the elements of the complex to communicate with each other (e.g. central processing complex to central processing complex, central processing complex to shared memory . . .). The present invention sets forth such a system for the transmission of information. Referring to the generalized data processing complex depicted in FIG. 1, there are physical links, 10-17, which interconnect the elements, E1-E5, of the complex. FIG. 1 is intended to depict a generalized complex of interconnected elements, and not any particular data processing complex. The elements E1-E5 can be, among other things, central processing complexes, I/O devices, Direct Access Storage Devices (DASD), or other electronic elements such as a commonly shared memory system. Each element of the data processing complex does not have to be interconnected with every other element (e.g. a printer does not have to be connected to a DASD). Any given data processing element can be connected to several other elements. The part of each element where interconnections are made to other elements is called a channel. As seen in FIG. 2, data processing element E3 has four different channels, 20-23, by which it is connected to four other elements, E1, E2, E4 and E5 by links 11, 13, 15 and 17 respectively.

Each of the physical links attached to a channel can contain a single conductor or be constructed from a number of conductors. In the present invention, the conductors of the link are preferably optical fibers, although the apparatus and method described here are equally applicable to a link constructed of copper conductors, radio waves, infrared, microwave or any other well known transmission medium.

The channel will also contain the necessary drivers and receivers which are appropriate for the particular transmission medium which is being employed. The optical fibers of the present invention are typically installed in pairs called fiber-pairs as seen in FIG. 3. In FIG. 3, Channels A and B are shown as being connected by a single conductor, fiber-pair 100 and 101. The designation of A and B for the two channels depicted in this figure is solely for explanatory purposes. Although there are two fibers, 100 and 101, this fiber-pair can be considered a single conductor because both fibers are required for bidirectional communication. Each fiber in the system interconnects a driver at one end and a receiver at the other end of the link. For example, fiber 100 connects transmitter 110 in Channel A with receiver 115 in Channel B (conversely fiber 101 connects driver 125 in Channel B with receiver 120 in Channel A). A driver and receiver are paired in a channel in a device called a transceiver. In FIG. 3, transmitter 110 and receiver 120 in Channel A form one transceiver, 130, while transmitter 125 and receiver 115 constitute a separate transceiver, 140, in Channel B. These interconnected transceivers allow for data to flow in one direction on one fiber (e.g. from Channel A to Channel B on fiber 100) and in the opposite direction on the other fiber (e.g. from Channel B to Channel A on fiber 101).

If there is only one set of transceivers interconnecting the channels, as depicted in FIG. 3, the information being transferred between devices would have to be transmitted in a strictly serial fashion. Having a single conductor will limit the bandwidth of information flow which can pass between elements of the data processing complex. The maximum bandwidth in a single conductor system will be governed primarily by the transmission medium chosen. In order to achieve enhanced performance via greater bandwidth, the present invention employs multiple transceivers and associated conductors in a parallel fashion as seen in FIG. 4. In FIG. 4, there are n conductors interconnecting Channels A and B, each of the conductors being a fiber-pair in the preferred embodiment. The n conductors collectively form a link or bus interconnecting the two channels. The information which is to be transferred across the link is broken up into smaller frames of information. Each of the frames is transmitted serially on a single conductor, but in parallel with other frames of information on the other conductors in the link. In this manner, the information can be transmitted serially on a single conductor in a parallel fashion on several conductors. This serial/parallel transmission of data provides a greater bandwidth and thereby enhanced performance of the communication system.

One of the difficulties in establishing a system in which there is parallel transmission of serial data over several conductors, is determining which of the conductors actually constitute the physical link between the connected elements. For example, in one channel of a computer element, there might only be four conductors attached to the channel, while in another channel in the same element there might be sixteen conductors. The system has to be able to identify and utilize the actual number of conductors attached to the specific channel. Furthermore, since it cannot be guaranteed that all of the conductors attached to a channel are operational when the system is initialized, a robust system cannot rely on a full complement of conductors. To be truly adaptable, the system has to be able to reconfigure the bus upon a failure, addition, deletion or repair of a conductor. This reconfiguration will employ a reduced number of conductors and a correspondingly reduced performance, but the loss of a conductor in the parallel bus will not render the entire bus inoperable. The present invention allows for all of

the above flexibility while maintaining the greatest bandwidth allowed by the channel implementations at each end of the link.

In the preferred embodiment of the present invention, there are 1, 2, 4, 8, 16, 32 or 64 transceivers in the link connecting any two channels. Each transceiver in the link uses one fiber to carry signals from one computer element to the other computer element and another fiber to carry signals in the opposite direction. It is possible for information to flow on both fibers at the same time. The speed and transmission characteristics of the transceivers may vary from technology to technology. However, all of the transceivers and conductors in a link will have the same general speed and transmission characteristics. For purposes of this discussion, when the term transceiver is used, it is ordinarily meant to imply the transceiver and associated fiber-pair to which it is connected.

In normal operation, the link will attempt to use all of the installed transceivers. At times, however, operation of the link has to proceed using fewer than the full complement of transceivers. For example, to enhance availability, facilitate maintenance, or accommodate the situation where the number of transceivers installed at one computer element differs from the number installed at the other connected element. The transceivers which actually participate in operational information transfer comprise what is called the operational link. The transceivers attached to the channel are assigned to the operational link either during a link-initialization or link-recovery process. Depending on the number of transceivers installed, the number of operational transceivers, the number of configurations supported by the particular channel at each end of the link, and the outcome of the initialization or recovery process, there are 1, 2, 4, 8, 16, 32 or 64 transceivers in the operational link. The normal operation of the link is to use the full complement of attached transceiver pairs. The present invention, though, allows operation of the link with fewer than a full complement of conductors.

At the beginning of the initialization of the link between two elements of the data processing complex, there are multiple conductors which connect the two elements as shown in FIG. 4, but it is unknown which of the conductors are available for operation. If the conductor is unavailable, it cannot form part of the operational link. The first step in the link-initialization procedure is therefore, to establish whether or not there is operational connectivity on each conductor which supposedly connects the two elements. It is an essential determination to initially ascertain which of the conductors are actually, physically capable of carrying a signal from one element of the data processing complex to the other element. In a system employing traditional copper conductors, this connectivity determination might be some type of information exchange test where a series of bits are transmitted across the link and then a response is transmitted back to the originating element to ensure connectivity. In more sophisticated transmission systems such as a radio waves or a fiber optic system, a connectivity test assumes a somewhat more complex nature in the form of a synchronization test. In these more complex systems, the information which is being transmitted is typically being converted from one form of energy to another and back again (e.g. from electrical signals to optical signals and back to electrical signals). Because of the complex nature of these conversions, a communication system must have more complex methods of ensuring connectivity between the terminating elements of the system. In an optical fiber system, connectivity is tested by a procedure known as synchronization. Synchronization involves two transceivers which are

connected by the fiber-pairs. During synchronization, a transmitter will attempt to send an encoded signal representing a bit stream to a receiver on the other end of the link. Whenever a signal is present on a transceiver, the receiver attempts to achieve synchronization on both bit transmission and transmission-unit boundaries of the received encoded bit stream. Bit synchronization is where the receiver attempts to identify the start of each bit of the received signal. A bit clock is thus established by this bit synchronization.

The other type of synchronization is known as transmission-unit synchronization. In this type of synchronization, the receiver must be able to identify the beginning and the end of a unit of data which is being transmitted by the driver on the other end of the link. A transmission-unit being a specified number of bits which forms a more manageable unit of data during inter-element communication. If the receiver is able to recognize the bit or unit boundaries, it is said to have achieved synchronization. In the preferred embodiment of the present invention, the system attempts to achieve unit synchronization on every transceiver attached to the channel. It does this by sending a signal from each of the transmitters in each of the transceivers. Once the receivers on the other side of the link have acquired unit synchronization, they inform their respective transceivers on the other side of the link of this synchronization-acquired state. After a specific period of time subsequent to the first transceiver attached to the channel having achieved synchronization (i.e. a timeout) or subsequent to all transceivers having achieved synchronization, the channel assumes that all the transceivers which are functional have established synchronization and every transceiver has achieved its final transceiver state. The final result of the connectivity, or more specifically synchronization, process is that the system has determined which of the physical conductors in the link are initially capable of becoming part of the final operational link between the two elements of the data processing complex.

The next step in the link initialization procedure is for the channel to invoke the initiative to acquire node descriptors from each of the functional transceivers attached to the channel. It must be noted that once any or all of the transceivers on each side of the link have achieved synchronization, the following steps in the process can be performed in parallel and independently in each of the two channels connected by the link. The system can be designed so that the synchronization and node descriptor process can be performed in a sequential fashion on each transceiver, or alternatively all of the transceivers can be synchronized and acquire node descriptors in parallel. A node descriptor is an identifier which uniquely describes a computer element, and optionally a particular interface of a computer element. Each element in the system will have a unique node descriptor and furthermore each interface of each computer element will have a unique node descriptor. The node descriptors for the interface on the other end of the link are acquired to ensure that every transceiver attached to the channel is also attached to the same location on the other end of the link. There are two methods of acquiring the node descriptor. In the first method, the node descriptors for the elements on the other end of the link are acquired by sending a command from each transceiver attached to the channel. This command is called an operational transceiver request. An operational transceiver request is sent on each of the transceivers independently. A transceiver which receives the request will respond back with its location's unique node descriptor. This response is called an operational transceiver response. Both sides of the link will originate and respond to operational

transceiver requests and responses. There is no timing relationship between the two sides during the operational transceiver request/response process (e.g. channel A on one side of the link can perform this operation in parallel and independently of channel B on the other side of the link). Both operational transceiver requests and operational transceiver responses may be sent and received at the same instant, or they may be sent in a sequential fashion.

In the second method, the node descriptors for the elements on the other end of the link are acquired by receiving a frame from each transceiver attached to the channel. This frame is called an operational transceiver frame. Operational transceiver frames consist of the node descriptor along with a request or response command. Only one end of the link must originate a request and the other end respond to exchange node descriptors. An operational transceiver request can be sent on each of the transceivers independently. A transceiver which receives the request will respond back with an operational transceiver response which contains its unique node descriptor. Both sides of the link can originate and respond to operational transceiver requests and responses, only one command must be received to acquire the node descriptor. There is no timing relationship between the two sides during the operational transceiver request/response process (e.g. channel A on one side of the link can perform this operation in parallel and independently of channel B on the other side of the link). Both operational transceiver requests and operational transceiver responses may be sent and received at the same instant, or they may be sent in a sequential fashion.

After a channel has received the first operational transceiver response or request at one of its transceivers, the channel verifies that the transceiver number in the node descriptor matches the transceiver number of the transceiver which received the response. For example, if transceiver number 5 received the first operational transceiver response, the channel would check to make sure that the received node descriptor indicated that the response came from transceiver number five on the other end of the link (i.e. transceiver number five is connected to transceiver number 5). If the transceiver numbers match in the first transceiver to receive an operational transceiver response, the channel places that transceiver in a transceiver-configured state. The node descriptor included in this first response is saved for comparison with later responses. With the receipt of every subsequent operational transceiver response or request, the subsequent node descriptors are compared to the node descriptor saved from the first transceiver which became configured. All subsequent transceivers must have the same node descriptor, with matching transceiver numbers, to enter the configured state. This checking ensures that all of the configured transceivers are actually connected to the same node descriptor (i.e. the same interface of the same computer element on the other side of the link). There is a time limit as to when a response from a transceiver must be received. This limit is implemented in order to prevent the system from hanging up while waiting for a transceiver which might never respond (e.g. a broken transceiver). There is also a retry operation in which the command to acquire the node descriptor is repeated a number of times if the previous command failed. There is a maximum number of retries which are attempted before the operation will fail. If the node descriptor comparison is successful and the transceiver-configured state is entered, the transceiver is marked as configured in a Configured-Transceiver table. If the transceiver returned an invalid node descriptor, or the limits for response time and retry have been exceeded, the trans-

ceiver is marked as being not configured in the Configured-Transceiver table. A sample Configured-Transceiver table is shown in FIG. 5.

The method for establishing the operational bus is primarily driven by state tables called Transceiver-State tables. Each channel of each element in the data processing complex contains a Transceiver-State table. This table actually contains two sub-tables, the Allowed-Operational-Link table and the Configured-Transceiver table. As described above, the Configured-transceiver table contains an entry for each of the transceivers attached to the channel. This entry will indicate whether or not the transceiver is in a configured state as can be seen in the sample Configured-Transceiver table of FIG. 5. If a particular transceiver was not operational in the sending channel, was considered off-line by the channel, the node descriptor did not match, or no operational transceiver response or request was received, the transceiver does not reach the configured state and is marked as not configured in the Configured-Transceiver table.

The next operation in the establishment of the operational link uses the Transceiver-State tables in order to determine which combination or set of transceivers is to be considered for the link. This operation results in an intermediate set of transceivers called an Intended-Operational-Link. The different combinations or sets of transceivers which could possibly become part of the Intended-Operational-Link are called Allowed-Operational-Links and reside in the Allowed-Operational-Link table, the second of the Transceiver-State tables. This table contains an ordered list of all of the allowed sets of transceivers which could possibly become an operational link for the channel. If a combination of transceivers is not in the Allowed-Operational-Link table, it can never become an operational link. It must be kept in mind that the Allowed-Operational-Link table does not indicate whether or not a transceiver or set of transceivers is actually operational, it only provides the various combinations of transceivers which could possibly form an operational link. A sample Allowed-Operational-Link table is depicted in FIG. 7 and will be discussed more fully later in this description. The Allowed-Operational-Link table is preferably built before the beginning of the link initialization, and would not be dynamically created such as the Configured-transceiver table. As an alternative to having an actual table in the memory of the channel. Each of the Allowed-Operational-Links could be dynamically generated at the point in the process at which they are required. The generating mechanism would output the permissible permutations of combinations of transceivers, governed by the same rules which govern the generation of the Allowed-Operational-Link table.

The preferred ordering of the sets of Allowed-Operational-Links in the Allowed-Operational-Link table is shown in FIG. 6. In examining FIG. 6, it can be seen that Allowed-Operational-Links are divided into successively smaller and smaller sets of transceivers. The first entry in the table will be a link containing all of the transceivers attached to the channel. The full set of transceivers attached to a channel is then divided into halves, then into quarters, then into eighths and so on until the smallest set of transceivers which form an Allowed-Operational-Link contains only a single transceiver. For example, to define a quarter of the transceivers, the full set of transceivers is divided into quarters starting with the highest order transceiver (transceiver 1, 2, . . .) for 1/4 of the transceivers in the link; then the second quarter of the transceivers starts, followed by the third and fourth quarters. This quartering divides an Allowed-Operational-Link into four sublinks that are each 1/4 of the size of the

original Allowed-Operational-Link and contain consecutively numbered transceivers in each link. Any given transceiver will only be contained in one of the quarter length sublinks. The process is the same for each of the sublinks; $+e, \text{fra } 1/2, 1/4, 1/8, 1/16, 1/32+ee$ and $1/64$ sublink. Each sublink contains $+e, \text{fra } 1/2, 1/4, 1/8, 1/16, 1/32+ee$ or $1/64$ of the total number of transceivers attached to the channel, starting from the highest order transceiver to the lowest and using consecutive transceiver numbering. It can be appreciated that the present invention is not limited by the 64 transceiver channel depicted in these examples. For example, if the channel happens to have 256 transceivers attached, then there will be granularity in the Allowed-Operational-Link table down to $1/256$ of the full set of transceivers.

FIG. 7 depicts a sample Allowed-Operational-Link table for a channel with a total of four transceivers attached to the channel. The sample table shown in FIG. 7 has been generated based on the methodology outlined in FIG. 6. The first entry in the sample table contains all four of the transceivers. The second entry in the table contains a set of transceivers including the first one half of the transceivers, namely the first and second transceivers. The third entry is the set containing the lower one half of the transceivers, the third and fourth. The fourth entry contains the first quarter of the transceivers, the first transceiver. The remainder of the table continues until the final allowed set of transceivers contains only the single transceiver, the fourth. It can be seen from an examination of both FIGS. 6 and 7 that only certain combinations, or sets, of transceivers can be considered when forming an operational link. For example, transceivers one, two, three and five could never form an operational link because they are not combined together as an entry in the Allowed-Operational-Link table. Furthermore, any combination of transceivers whose total number is not a power of two is not allowed to form an operational link. The reason for the methodology of allowing certain combinations to form an operational link and disallowing other combinations is based on the capabilities designed into the channel. A person experienced in the art can see that any combination of transceivers is theoretically possible, but implementation considerations may make some of the combinations less desirable. For example if every single possible combination of transceivers were allowed, the channel would have to incorporate a full crosspoint switch which would have an effect on the performance of the system. The use of the Allowed-Operational-Link table permits the designer to implement only a subset of all possible combinations of transceivers, which results in lower cost implementations.

Returning to the procedure for determining the Intended-Operational-Link, once the Configured-Transceiver table has been built, and all of the transceivers have completed their respective entries in the table, or a specific period of time has expired, the next step in the link initialization procedure is invoked if one or more transceivers is in a configured state. If no transceivers were able to be configured, or fewer than a minimum number, no operational link can be formed and link initialization fails. If there are a minimum number of configured transceivers, the Allowed-Operational-Links are systematically compared to the transceivers which have been configured, until a match is found or until a minimum number of transceivers cannot be found for the Intended-Operational-Link. This process begins by searching the Allowed-Operational-Link table for the different sets of transceivers which comprise Allowed-Operational-Links. The search will first result in the Allowed-Operational-Link which is the set containing the full

complement transceivers. Once the search has acquired this set, the system then checks the Configured-Transceiver table to see if each of the transceivers in the set has been configured. If all of the transceivers in this first Allowed-Operational-Link are configured, the procedure will move on to the next step described below. If even one of the transceivers is not configured (as indicated in the Configured-Transceiver table), then the comparison will fail and the system will continue to search the Allowed-Operational-Link table for the next set of transceivers in an Allowed-Operational-Link. A second search of the Allowed-Operational-Link table will result in the set of transceivers in the high order one half of the transceivers (the second entry in the Allowed-Operational-Link table). The Configured-Transceiver table will then be consulted again to determine if each member of the second set of transceivers has been configured. The process of searching and comparing continues until a match is found, the Allowed-Operational-Link table is exhausted, or a minimum number of transceivers for the operational link is not satisfied. A match is found when every transceiver in an Allowed-Operational-Link is found to be configured in the Configured-Transceiver table. Once a match is found, this set of transceivers becomes an Intended-Operational-Link. If the system cannot establish an Intended-Operational-Link, then link initialization will fail. As a reminder, each of these process steps can be performed independently and in parallel in each of the two channels on either side of the link. Therefore, at some point, there can be two Intended-Operational-Links defined, one in each channel on either side of the link.

As explained previously, as an alternative to the Allowed-Operational-Link table, the system can dynamically generate the Allowed-Operational-Links as required. The first Allowed-Operational-Link generated would be the full set of transceivers attached to the channel. If this set contained a transceiver which was found not to be configured (in the Configured-Transceiver table), then the system would dynamically generate the next Allowed-Operational-Link, the high order one half of the transceivers. The process would continue similarly to the one using the table, until either an Intended-Operational-Link is formed, or link initialization fails. The choice of whether to use a table or to dynamically generate the Allowed-Operational-Links, entirely depends entirely upon the particular implementation. If there are only a small number of Allowed-Operational-Links, then it might be preferable to dynamically generate the small list. Alternatively, if memory space in the channel is not limited, then a large Allowed-Operational-Link table could reside in the memory hardware which might add additional speed to the process of establishing an Intended-Operational-Link.

Once the Intended-Operational-Link has been established, the system can proceed onto the next step in the process of determining the final operational link. This step involves ensuring that both sides of the link agree on which of the conductors between the two computer elements will form the operational link. Up to this point, each channel can independently determine its own Intended-Operational-Link, but there has been no verification that the two Intended-Operational-Links correspond. At this point in the procedure, there are three methods by which the system can ensure that both sides agree on the operational link. The first method involves initiation on both sides of the link in a peer-to-peer relationship, while the second and third methods employ control by one channel in a master/slave relationship. The first and second methods require the use of a command, called an establish-transceiver-group command

13

(ETG command), which is sent on each transceiver in the Intended-Operational-Link. The third method requires the use of a command, called an establish-link command (EL command), which is sent on each transceiver in the transceiver-configured state.

In the first method of verifying the Intended-Operational-Link, each channel issues an ETG command on each transceiver in its respective Intended-Operational-Link. This command will include a designation of each of the fibers which is a member of the Intended-Operational-Link. As the channel on the other end of the link begins to receive the ETG commands on its transceivers, it will compare them to its own Intended-Operational-Link in order to verify that there is a correspondence. The receiving channel will test to see that each of its transceivers which received the ETG command are included in its own Intended-Operational-Link. If there is one-to-one correspondence, then the channel which received the ETG commands will send back a response indicating this state, on each of the transceivers in the verified Intended-Operational-Link. This response is called an ETG response. Both ends of the link issue the ETG command on each transceiver in its Intended-Operational-Link, and the ETG response must be received back on the same set of transceivers. If a channel receives responses on all of the transceivers in its Intended-Operational-Link, then the Intended-Operational-Link becomes the operational link. If the ETG commands are received on a set of transceivers which don't agree with the channel's own Intended-Operational-Link (too few or too many transceivers), then the channel will send back a failed-transceiver group response and the initiating channel will attempt to resend the ETG commands up to a maximum number of times. If the channel issues ETG commands the maximum number of times without responses or failed responses, then link initialization will fail. This first method of verifying the Intended-Operational-Link can be seen to be a peer-to-peer operation in that both channels on either side of the link issue commands and formulate responses. Neither of the channels is acting in a dominant role.

In the second method of verifying the Intended-Operational-Link, only one channel, a master, originates ETG commands on the transceivers in its Intended-Operational-Link. The channel on the other side of the link, the receiving or slave channel, does not originate any commands. The slave channel will observe which of its transceivers receive ETG commands in order to verify that they comprise a fully configured and Allowed-Operational-Link. In order to perform this verification, the receiving channel uses its Configured-Transceiver table and Allowed-Operational-Link table. If the Intended-Operational-Link is, for any reason, unacceptable to the slave channel, it issues a failed-transceiver-group response back to the master channel. As with the peer-to-peer method of verifying the Intended-Operational-Link, if the master receives a failed-transceiver-group response, it will attempt to send the ETG command up to a maximum number of times. If the master issues the ETG command the maximum number of times with no responses or failed responses, then link initialization will fail. If the set of transceivers which received an ETG command comprises a fully configured, allowed operational link, the receiving channel sends ETG responses indicating as such on all of the transceivers in the verified Intended-Operational-Link. The master channel examines the responses to ensure that responses were received on all of the transceivers of its Intended-Operational-Link, and no other transceivers. If for any reason the responses do not match its Intended-Operational-Link, the intended link does not become the opera-

14

tional link. If there is one-to-one correspondence between the Intended-Operational-Link and the transceivers which received ETG responses, then the Intended-Operational-Link becomes the actual operational link. In this second method, only one end of a link issues commands on the transceivers in its Intended-Operational-Link, and the responses must be received on the same set of transceivers. Regardless of which method (peer-to-peer or master/slave) is used to verify the Intended-Operational-Link, the result of this pad of the process is that an operational link has been established.

In the third method of verifying the Intended-Operational-Link, only one channel, a master, originates EL request commands on all the transceivers in the transceiver-configured state. A field in the EL request command indicates if the transceiver is capable of being part of an operational link at the master end of the link. The channel on the other side of the link, the receiving or slave channel, does not originate any requests. The channel which receives this command, a slave, will examine the frames and determine the set of transceivers which can become an operational-link. A logical AND operation is performed on a transceiver basis for the state of the transceiver at the slave channel as determined by the transceiver-configured table entry and the capability of the transceiver to become a member of an intended-operational-link as indicated in the frame received by the slave on the respective transceiver. The results of this logical operation is performed for each of the transceivers and the results used as the candidate transceivers in a search of the allowed-operational-link table and a response is sent from the slave to the master indicating the results on each transceiver in the transceiver-configured state. The response frames returned contain a bit that is unique on each transceiver frame that indicates that the transceiver is pad of the intended-operational-link at the slave. If no Intended-Operational-Link is possible, it issues a EL response command with no intended-operational link. Once the slave sends these response frames, the intended-operational link becomes the operational link at the slave. Once the master receives these responses, the intended-operational-link as defined by the received frames is checked against the allowed-operational-link table entries to ensure that the intended-operational-link is valid, and if it is valid, the intended-operational-link becomes the operational-link and use of the link can begin. The master must receive responses on all transceivers in the transceiver configured state before the intended-operational-link can be checked. As with the peer-to-peer method of verifying the Intended-Operational-Link, if the master receives an EL response command with no intended-operational-link specified, it will attempt to send the EL request command up to a maximum number of times. If the master issues the EL request command the maximum number of times with no responses or failed responses, then link initialization will fail. If responses are not received within a specified time limit, the requests are resent. The allowed-operational-link tables are the same at both ends of the link. In this third method, only one end of a link issues requests on the transceivers in the transceiver-configured state.

Regardless of which method is used to verify the Intended-Operational-Link, the result of this pad of the process is that an operational link has been established.

For the first and second methods, before message operations using the operational link can begin, a command called a bundle request is issued. The purpose of this command is a final assurance that both channels on either side of the link are in agreement as to which fiber pairs form the operational link. This command will only be sent once a channel has

received establish-transceiver-group responses on all of the members of its Intended-Operational-Link and the operational link has been determined. If the channel receiving a bundle request agrees on the operational link, it sends back a bundle response and operations can commence on the operational link. If, for some reason, the receiving channel does not agree on the operational link, it will not send back a bundle response and no operational link is established. The channel sending a bundle request will wait a threshold time period within which it is expecting a bundle response. If no response is received during this threshold, the channel will reissue the bundle request. The bundle request can be reissued up to a maximum number of times. The threshold time for a bundle response and the maximum number of reissues of the request can be adjusted by the designer of the system to the requirements of particular application. For example, in the preferred embodiment of the present invention, the time threshold is set to one second for a bundle response and the bundle request can be reissued a maximum of four times.

It is possible that a bundle response or EL response will be issued by one channel and never received by the other channel. In this situation, the channel which issued the bundle response or EL response will believe that there is an operational link in place and the other channel does not (it is still waiting to receive the bundle response or EL response). If a channel which is expecting a bundle response or EL response receives another valid message on the link, it will assume that the bundle response or EL response with the accept field set on was sent by the other channel and will configure itself accordingly. This means that it will establish the operational link from the Intended-Operational-Link and full communication on the link can proceed as if the bundle response or EL response was never lost.

After the operational link is fully established, the next and final operation in the initialization of the communication link between the two computer elements is to set up message buffers in the connected channels. The basic structure of the message buffers employed in the present invention is depicted in FIG. 9. A message buffer comprises two separate buffer structures which are required to support the transmission of a single message. These buffer structures are called the originator buffer, 500, and the recipient buffer, 550. A message buffer performs the task of buffering the information which is to be transmitted across the operational link. For example, if Channel A is initiating a message across the link, its originator buffer, 500, is used to hold the information before transmittal across the link. The originator buffer acts as a staging area for the message. As explained previously, information is transmitted across the link in a serial/parallel fashion over several conductors. The information is taken from the originator buffer, broken up into the pieces called frames, which are then transmitted in a serial manner on the parallel fibers. The recipient buffer, 550, in receiving Channel B, performs the converse function of reassembling the information comprising the message when the individual frames are received at the channel. The breaking up and reassembly of information for transmission is sometimes called interleaving. Although the buffer, 550, in Channel B is called a recipient buffer, it can, and does originate communications back to Channel A. For example, an ETG response from Channel B would be sent from its recipient buffer 550 to the originator buffer, 500, in Channel A. In this example, the originator buffer, 500, in Channel A would be the area in which the ETG response was reassembled after transmission across the link.

Both the originator and recipient buffers have three logical areas: a request area, 501 and 551; a response area, 502

and 552; and a data area, 503 and 553. The requested information for a particular message is sent from the request area, 501 of the originator buffer, 500, to the request area, 551 of the recipient buffer 550. Similarly, any response information is sent from the response area, 552, of the recipient buffer, 550, to the response area, 502 of the originator buffer, 500. The data areas, 503 and 553, of each of the buffers are used for transmission and receipt of any additional data which is required for requests and responses.

A single paired set of originator and recipient buffers, called a message buffer, is required for the transmission of a single message. In the present invention, channels are capable of concurrent transmission of several messages. The buffer structure required for this type of communication is depicted in FIG. 10. It can be seen in the upper half of FIG. 10 that there are $n+1$ originator buffers in Channel A and a corresponding $n+1$ number of recipient buffers in Channel B. These $n+1$ paired buffer structures can support $n+1$ simultaneous message operations. For each originator buffer in Channel A, there exists a paired recipient buffer in Channel B. For example, originator buffer 600 in Channel A is paired with recipient buffer 610 in Channel B and this pair will support a single message operation. The originator buffer marked number 1 in Channel A is paired with the recipient buffer marked number 1 in Channel B. This second pair of buffers will support a second, concurrent message operation. The example in FIG. 10 depicts a system in which there are an equal number of recipient buffers and originator buffers, but there is no requirement that there be a originator buffer for every recipient buffer. The only requirement is that there is a recipient buffer for every originator buffer in the system. For example, Channel A in FIG. 10 might have only two originator buffers and Channel B could have 10 recipient buffers. In this example, the two originator buffers in Channel A would be paired with two of the recipient buffers in Channel B. The rest of the recipient buffers in Channel B would be unpaired and would not support message operations. In this example, the system would only be able to support two concurrent message operations from Channel A to Channel B.

It can be seen from further examination of the lower half of FIG. 10 that messages can be initiated from either channel. This portion of the Figure illustrates that channels are capable of containing both originator and recipient buffers. If Channel B was to initiate a message operation, it would use its originator buffer, 620, to send the message to Channel A's recipient buffer, 630. The number of recipient buffers in a channel does not have to equal the number of originator buffers in that same channel. In Channel B depicted in FIG. 10, there are $n+1$ recipient buffers, while there are $m+1$ originator buffers. M does not have to be equal to, greater than or less than n . The only requirement is that each originator buffer in a channel is paired with a recipient buffer in the channel on the other end of the link. In the example depicted in FIG. 10, there are $n+1$ recipient buffers in Channel B paired with $n+1$ originator buffers in Channel A. Similarly there are $m+1$ originator buffers in Channel B, paired with $m+1$ recipient buffers in Channel A.

The example depicted in the lower half FIG. 10 shows $m+1$ originator buffers paired with $m+1$ recipient buffers. As explained previously though, the only requirement is that there be one recipient buffer for each originator buffer. For example, Channel B might have four originator buffers and Channel A could have the capability to support seven recipient buffers. The four originator buffers in Channel B would be paired with four recipient buffers in Channel A and the remaining three recipient buffers in Channel A would not

be utilized. The type, number, and size of the buffers on either end of the link is limited by the memory available to the channel and the capabilities built into the particular channel. Therefore, the $n+1$ number of originator buffers in Channel A, and the $m+1$ number of originator buffers in Channel B is the maximum number of originator buffers which could be established for the system depicted in FIG. 10. There does not have to be a one to one correspondence between the number of originator and recipient buffers, but there has to be one recipient buffer for every originator buffer in order to establish a message buffer.

Because of the requirement that each originator buffer must be paired with a corresponding recipient buffer on the other side of the link, both channels must agree and participate in the establishment of the type, number and size of the buffers which are to be used for message operations. In order to establish the type, number and size of the buffers which are to be used, special commands and responses called buffer commands and responses are exchanged between interconnected channels. Each channel will inform its counterpart channel on the other end of the link, of its capacity for recipient buffers using a set-buffer-size request. For the example depicted in FIG. 10, Channel B would inform Channel A that it is capable of providing $n+1$ recipient buffers of a specified size (size of request and data areas) to support message operations. Channel A, upon receipt of the set-buffer-size request evaluates whether or not the $n+1$ recipient buffers of the specified size are sufficient to support message operations. If the recipient buffers are insufficient to support message operations, Channel A will not establish any originator buffers, and message operations will not commence. If the $n+1$ recipient buffers are satisfactory to Channel A, it will establish the number of originator buffers which it requires for message operations. As stated previously, the number of originator buffers which Channel A establishes does not have to equal the number of recipient buffers capable of being provided by Channel B. For example, if Channel B indicated in its set-buffer-size request that it was able to provide 10 recipient buffers, but Channel A only required 2 originator buffers for message operations, then Channel A would only establish 2 originator buffers and only two message buffers would be established.

For each message buffer that exists (i.e., originator/recipient buffer pair), one message operation can take place. Thus, the multiple message buffer structure illustrated in FIG. 10 allows multiple messages to be in execution concurrently. Furthermore, these message operations can occur simultaneously in both directions on the link. Messages can originate from either Channel A or Channel B.

A channel issuing a set-buffer-sizes request will wait a threshold period of time during which it will expect a response. If no response is received within this threshold, the assumption is that the request was somehow lost on the link, the request times out and another set-buffer-sizes request is sent. The set-buffer-sizes request can be resent up to a maximum number of times. The same request is sent every time (i.e., the number and size of the buffers contained in the request does not change). If the maximum number of retries is reached, then link initialization will fail.

If the channel which receives a set-buffer-sizes request concurs in the number and sizes of the buffers (they meet or exceed the requirements for message operations), then it sends back a set-buffer-sizes response. Both channels will initiate set-buffer-sizes requests and will respond with set-buffer-sizes responses. If a channel does not have the capability of supporting recipient buffers, it still initiates a set-buffer-sizes request, but the number of buffers included

in the request will be zero. Once each channel has initiated both request and responses, the link is fully established and operational, and true message operations call commence.

A second method for establishing the buffer sizes and number uses a single request/response exchange in a master/slave relationship. The master end of the link sends a set buffer sizes request with the number and sizes of the recipient buffers at the master end of the link. When the slave end of the link receives the set-buffer-sizes request from the master end of the link it verifies that the required number and sizes of recipient message buffers exist at the master end of the link and sends set-buffer-sizes response that indicates the slave's recipient message buffer number and sizes, and a field that indicates the slave end's acceptance (or rejection) of the master's recipient message buffers. If the master receives a set-buffer-sizes response that indicates the slave's acceptance of the message buffers and the slave's recipient message buffer number and sizes are acceptable to the master, then a bundle request is sent and the description for the bundle request is used. Once the bundle request/response exchange is complete, message exchanges can begin. If the master receives a set-buffer-sizes response indicating the slave's non-acceptance or the slave's message buffer number or sizes are not acceptable to the master, then the link-initialization procedure fails.

A master channel issuing a set-buffer-sizes request will wait a threshold period of time during which it will expect a response. If no response is received within this threshold, the assumption is that the request was somehow lost on the link, the request times out and another set-buffer-sizes request is sent. The set-buffer-sizes request can be resent up to a maximum number of times. The same request is sent every time (i.e., the number and size of the buffers contained in the request does not change). If the maximum number of retries is reached, then link initialization will fail.

The above methods of establishing an operational parallel link were primarily directed at initially establishing an operational bus when the system of computer elements or a new link is first configured. One of the strengths of this method, though, arises when one of the conductors of the parallel bus becomes disabled or otherwise must be taken offline, and the bus must be reconfigured. In prior communication systems, the loss of one of the conductors of a parallel bus usually meant that all intersystem communication would cease until the conductor was either fixed or replaced. With the present invention, a link recovery procedure is employed to reconfigure the link, excluding the malfunctioning conductor, and communication using the remaining functional conductors can resume. Operations will proceed with a degraded performance, but it is not necessary to halt communications entirely.

When one or more of the transceivers in the operational link enters a failed state or must otherwise be taken offline (e.g. for maintenance), the link recovery procedure is invoked. A transceiver will enter a failed state for several reasons including: removal of a conductor; failure of a transmitter or receiver; discontinuity in the conductor; loss of power; improper cable length causing signal degradation; or excessive splices in the cable causing signal degradation.

There are two methods of recovering an operational link. The more drastic and disruptive method is to begin the link initialization procedure described above. This method involves going through each of the method steps which have been fully described above, including: synchronization; node descriptor exchange; building of a Configured-Transceiver table; establishment of an Intended-Operational-Link

and so on. The less disruptive method is called a link-recovery procedure and only involves a subset of the steps of the full link initialization procedure.

When the system enters the link-recovery procedure, the Configured-Transceiver table is updated in order to indicate the transceiver or transceivers which are now unavailable for use by the system. The system will then go onto establishing a new Intended-Operational-Link. The establishment of this new Intended-Operational-Link differs from the link initialization procedure in one significant area. Instead of searching the Allowed-Operational-Link table from the beginning, the system starts the search from the point where the last search ended. This is to say, the system knows that the new Intended-Operational-Link can only be a subset of the last operational link (i.e. it has to have fewer transceivers because at least one of the transceivers in the last operational link has failed). The last operational link was established from the last Intended-Operational-Link, therefore the search for the new Intended-Operational-Link can start in the Allowed-Operational-Link table from the point where the last Intended-Operational-Link was found. If the Allowed-Operational-Links are being generated by a software program, then the software program can be designed to begin generating Allowed-Operational-Links that are a subset of the last operational link (i.e. the last Allowed-Operational-Link which the software program generated).

Once a new Allowed-Operational-Link is either found or generated, the link-recovery procedure establishes the Intended-Operational-Link in the same manner as the link initialization procedure. It begins by testing the set of transceivers against the Configured-Transceiver table, determining if the set forms a fully configured operational link. The search and test steps are continued until a match is found. Once the match is found, that set of transceivers becomes the new Intended-Operational-Link. As in the link initialization procedure, the link-recovery procedure sends out a Establish-Transceiver-Group, ETG, or establish-link, EL, command on each transceiver in the new Intended-Operational-Link. The message exchange involved in the ETG or EL process for link recovery is the same as the ETG or EL process for link initialization, namely: the Intended-Operational-Link is verified in either the master/slave or peer-to-peer method; ETG or EL responses are received on each transceiver in the Intended-Operational-Link; the operational link is established; and the bundle request is optionally issued. The link-recovery procedure differs from the link initialization procedure at this point, in that no set-buffer-sizes commands need to be issued in link recovery. The buffer structures which were established during initialization are not disturbed during recovery. The system will utilize the same buffers, with a reduced number of transceivers in the link. While the system is in the process of link-recovery, all requests for message transmissions are presented with link busy indications. Message operations will not continue until the new operational link has been established.

FIG. 8a depicts a system according to the present invention in which all four fiber-pairs of a four fiber-pair link are operational. The message areas, 200 and 210, correspond to the message buffers described in relation to FIG. 9 and 10. These areas serve as buffers for incoming or outgoing messages. The interleave logic, 220 and 230, is the mechanism by which the information constituting messages is either broken up into frames for transmission or reassembled back into the original message after reception. In the case of interleave logic 220, the logic divides any message into four frames of equal size before transmission. This is the mean-

ing of the symbol 1:4 in the interleave logic box, 220. The interleave logic, 220, in FIG. 8a divides the information into four equal frames because the operational link in this particular example is made up of four fiber-pairs. During transmission of the message, each of the fiber-pairs will carry one of the frames of information in parallel with the other three fiber-pairs. Symmetrically, the interleave logic, 230, on the other side of the link reassembles the data back into the original message at a 4:1 ratio. The interleave logic, 230, receives four separate frames of information and proceeds to interleave the four frames back into the original message.

FIG. 8b(1) illustrates the four fiber-pair system of FIG. 8a in which either fiber-pair 3 or fiber-pair 4 are inoperative. Recall that the terms fiber-pair and transceiver can be used interchangeably when discussing the conductors in a link. This figure can be depicts a system which has just been initialized and the fiber-pair was initially inoperative or, alternatively, this figure depicts a system in which a fiber-pair became inoperative during message operations and the system has gone through the recovery procedure described earlier. In either case, the operational link in FIG. 8b(1) consists of only fiber-pairs 1 and 2. In this system, it does not matter if fiber-pair 3, fiber-pair 4 or both fiber-pairs are inoperative. If either one, or both are inoperative, then neither of the fiber-pairs can take part in an operational link. At this point it will be useful to explain this example by illustrating the method of determining an Intended-Operational-Link as has been described earlier.

The first step after synchronization has occurred is to build the Configured-Transceiver table. In this example, the entry for either fiber-pair 3 or 4 or both fiber-pairs, indicates that the fiber-pair(s) is NOT CONFIGURED (see FIG. 5). This entry can occur, for example, because the system was not able to achieve synchronization on the fiber-pair(s), an invalid node descriptor was returned on the fiber-pair(s) or no operational transceiver response was ever received after successive operational transceiver requests. After the Configured-Transceiver table has been built, the next operation is to search the Allowed-Operational-Link table for a link which is allowed to become an operational link. The Allowed-Operational-Link tables for this particular example will be the same as the one depicted in FIG. 7. (Remember, there are two Allowed-Operational-Link tables in the system, one in each channel attached to the particular link.) The first Allowed-Operational-Link in the table will be a set of transceivers containing all of the fiber-pairs, 1, 2, 3 and 4. The system will compare the fiber-pairs in this set to the status of each transceiver in the Configured-Transceiver table. During this comparison, it is found that this link is unacceptable since fiber-pair 3 or 4 or both fibers are NOT CONFIGURED and therefore cannot become part of an operational link. The second search of the Allowed-Operational-Link table will result in a set of transceivers including fiber-pairs 1 and 2. The testing of these transceivers against the Configured-Transceiver table, reveals that this is an acceptable link and fiber pairs 1 and 2 become the Intended-Operational-Link. The remainder of the procedure continues until fiber-pairs 1 and 2 are finally configured as the operational link. The important thing to identify at this point, though, is that it did not matter whether it was fiber-pair 3, 4 or both which were inoperative, the process of the preferred embodiment of the present invention of searching the Allowed-Operational-Link table and testing the Configured-Transceiver table would have driven both fiber-pairs 3 and 4 out of the operational link.

Turning to FIG. 8b(2), this figure depicts a system in which fiber-pairs 3 and 4 form the operational link. As with

the previous example, it does not matter if fiber-pair 1, 2 or both were inoperative, the method of the preferred embodiment of the present invention will establish an operational link out of fiber-pairs 3 and 4. At the point in link initialization, where the Allowed-Operational-Link table is being searched, the first set of transceivers found, which includes all of the transceivers, is unacceptable since either fiber-pair 1, 2 or both are NOT CONFIGURED in the Configured-Transceiver table. The second set of transceivers is similarly unacceptable since it contains only fiber-pairs 1 and 2. The third Allowed-Operational-Link is acceptable since it contains only CONFIGURED transceivers, namely fiber-pairs 3 and 4. Recall, as an alternative to the Allowed-Operational-Link table, the Allowed-Operational-Links can be generated by a software program. Using either the table or the software program will result in the same operational link, because both are constrained by the rules illustrated in FIG. 6. In both FIGS. 8b(1-2), the interleave logic, 220 and 230, is described as being 1:2 (or respectively 2:1). The reason for this is that since there are only two active conductors in the operational link, the information which is to be transmitted across the link has to be partitioned into two separate frames of data (as opposed to four frames in the prior example in FIG. 8a).

FIGS. 8c(1-4) depict systems in which there is only one active fiber-pair in the operational link. Looking at the Allowed-Operational-Link table in FIG. 7, it can be seen that each individual fiber-pair can form an operational link by itself. A system using only one conductor as the link between computer elements will result in the lowest bandwidth of any link, but the two computer elements will still be able to communicate on the one conductor. In these systems there is essentially no interleaving since the data is to be transmitted in a strictly serial fashion over the single conductor (hence the interleave ratio of 1:1).

In order to illustrate the recovery procedure described above, let us suppose that the current operational link is the one depicted in FIG. 8b(2), with both fiber-pairs 1 and 2 being inoperative. If fiber-pair 3 is somehow broken, the system will enter the recovery procedure. The first step in the recovery procedure is to mark fiber-pair 3 as NOT CONFIGURED in the Configured-Transceiver table. The next step will be to search the Allowed-Operational-Link table for a new acceptable link. The last acceptable Allowed-Operational-Link included fiber-pairs 3 and 4. The search of the table during the recovery procedure will begin at the next entry below this point since all of the above sets of links are known to contain NOT CONFIGURED transceivers (see FIGS. 5 and 7). The next three entries in the table yields sets of transceivers including fiber-pairs 1, 2 and 3 all by themselves. All three of these sets will be rejected during the testing operation since all three are NOT CONFIGURED. The final Allowed-Operational-Link is fiber-pair four all by itself. It constitutes an acceptable link because it is seen to be CONFIGURED in the Configured-Transceiver table. The remainder of the process continues as described above, e.g. using the ETG requests and responses, using the bundle requests and responses but not using the set-buffer-sizes (because the buffers remain the same regardless of how many conductors are in the operational link). The final result of the link-recovery procedure is the system depicted in FIG. 8c(4) with the operational link being made up of only a single fiber-pair, fiber-pair 4.

In order for a conductor to operate at its optimal bandwidth, the speed of the system logic at each end of the conductor must be taken into account. In many instances the system logic at one end, or both ends of a conductor may not

be able to keep up with the conductor transmission rate. In order to achieve the optimal utilization of the conductor's bandwidth, the effective bandwidth of the conductor can be slowed down to account for the system logic clock speed at each end of the conductor. A stuffing symbol can be inserted at the sending end and removed at the receiving end which allows the effective bandwidth of the conductor to be slowed down to that of the systems at each end of the conductor, while the conductor transmits at its normal bandwidth. The bandwidth of the receiver is communicated to the sender, allowing the sender to insert enough stuffing symbols that the sender's and receiver's effective bandwidth are both met. This allows optimal utilization of a conductor where the system clocks at each end can be different.

In a system that uses serial data transmission with an encoding similar to the 8/10 code described in U.S. Pat. No. 4,486,739 of Franaszek and Widmer entitled "Byte Oriented DC Balanced (0,4) 8B/10B Partitioned Block Transmission Code" (Franaszek and Widmer 8/10 code), and the unit of transmission over the serial link is composed of four characters, the stuffing symbol is called a null word. The following items are considered in the null word design:

1. The null word should have a high transition density. Since it may be repeated every other word, it should still not impede the phase lock loop from acquiring bit synchronization.

2. No words preceding and following the null word (including the null word) can cause false word-synchronization detection (single occurrence).

3. The null word should be a modification of the idle word. This property allows similar decoding by any hardware of both the idle and null words.

4. The null word should have zero disparity. Having zero disparity allows insertion and deletion in the data synchronizers at the system clock to conductor clock interface, which does not include the 8/10 encoder/decoder hardware.

5. The properties above make K28.5-D21.5-D21.5-D16.5 an acceptable choice for the null word. It is a modified idle word (K28.5-D21.5-D21.5-D21.5) with the last D21.5 changed to a D16.5. This null word has two codes (1100000101 1010101010 1010101010 0110111010 and 0011111010 1010101010 1010101010 1001001010), which are selected based on the running disparity of the data stream. Transmitters must use the proper version of the null word so that all of the 8/10 encoding rules are met. Reception should not require checking of the null encoding choice based on running disparity because of its limited value and costly implementation.

Transmission and reception of the null word have the following rules:

1. Null-word transmission and reception is independent of any frames that are sent on the conductor. Null words are allowed anywhere.
2. All transceivers on a channel must have the same null pacing value. Pacing is described later.
3. Each transceiver can send nulls independently.
4. Send skew rules must allow for inserted null words as part of the random skew (RSKEW) for the operational link. Random skew is explained in the continuation-in-part of U.S. patent application Ser. No. 07/839,986 filed Feb. 20, 1992 entitled "Frame Group Transmission and Reception for Parallel/Serial Buses" (IBM Docket No. PO9-92-001) now U.S. Ser. No. 08/071,146, filed Jun. 1, 1993 which is considered to be incorporated herein.

The null pacing parameter is a multiple bit value, which is the maximum number of non-null words between null words. Eight bits being a representative size for this value.

The transceivers send this parameter during the transceiver initialization process along with the node descriptor. The transceiver computes this value based on the ratio of its system clock rate to the link clock rate. The slower the system clock, the smaller the parameter is. When the value is '00'x, the transceiver is indicating to the other end of the link that no null words need be inserted by the other end of the link in its transmitted data. When the value is '01'x, the transceiver requires every other word in its received data to be a null word.

When a link is being initialized, various Continuous Sequences are exchanged. These sequences are word extensions of the ESCON sequences.

At this time, the transceivers have not yet received a null pacing parameter, and they transmit with a null pacing parameter of '01'x. This maximum null word insertion rate slows down the link bandwidth by 50%, which is the slowest a transceiver is allowed to run in order to operate. The transceivers continue to transmit with this null pacing value until the null pacing parameter is received from the other end of the link. Depending on the system clock speed at the other end of the link, the received pacing parameter may be either larger or smaller than the transmitted parameter. If the received parameter is larger, the other end of the link can operate faster than the transceiver; and if the parameter is smaller, the other end of the link operates slower than the transceiver. In either case, the link bandwidth has to be slow enough to accommodate both ends. So, the transceiver then compares its transmitted null pacing parameter with the parameter received from the other end of the link. The transceiver selects the smaller of the two parameters (except '00'x) and changes its transmit null pacing from '01'x to this new value. With this procedure, the transceiver is sending data at a rate which does not overrun its transmit capabilities, and receives data at a rate which does not overrun its receive capabilities.

When a transceiver detects a link failure, it sends a Continuous Sequence. At this time, the transceiver uses a null pacing parameter of '01'x since the physical link connection may have changed, and the transmitter is not sure of the capabilities of the other end of the link. The null pacing parameter will be reestablished during the transceiver initialization process.

Null words may be inserted by the sender at a rate higher than the established null pacing parameter. The rate must be limited to prevent the message from timing out.

FIG. 11 illustrates a receiver circuit which deletes the null words as they are received. The receive hardware consists of a RECEIVE REG (1100) which is 40 bits wide, a NULL DETECT circuit (1101) which examines the data for the NULL WORD, and an ASYNCHRONOUS BUFFER (1102) which is similar to the one described in U.S. Pat. No. 5,003,558 of Gregg entitled "Data Synchronizing Buffer for Data Processing Channels". The RECEIVE DATA (1103) and the RECEIVE CLOCK (1104) come from the deserializer, and the SYSTEM CLOCK (1105) comes from the channel hardware. The SYSTEM CLOCK (1105) may run either faster or slower than the RECEIVE CLOCK (1104). When the SYSTEM CLOCK (1105) runs slower than the RECEIVE CLOCK (1104), sufficient NULL WORDS are in the data stream to prevent an OVERRUN condition in the ASYNCHRONOUS BUFFER (1102). The outputs all go to the channel. They consist of the SYNCHRONIZED DATA (1106), a DATA AVAILABLE line (1107) which is used to signal valid data at the output of the ASYNCHRONOUS BUFFER (1102), and an OVERRUN

line (1108) which signals that a word of data in the ASYNCHRONOUS BUFFER (1102) was discarded.

The data on the RECEIVE DATA lines (1103) is clocked into the RECEIVE REG (1100) by the RECEIVE CLOCK (1104). If the data is a NULL WORD, the loading of the ASYNCHRONOUS BUFFER (1102) is suppressed by the AND circuit (1109). When the SYSTEM CLOCK (1105) is slower than the RECEIVE CLOCK (1104), NULL WORDS are required in the data stream. Transmission errors in the RECEIVE DATA (1103) may alter the NULL WORD so that it is not detected by the NULL DETECT circuit (1101). Such a condition may cause an OVERRUN in the ASYNCHRONOUS BUFFER (1102). This OVERRUN condition is handled by the ASYNCHRONOUS BUFFER (1102) by discarding a word of data. The OVERRUN line (1108) is used by the ASYNCHRONOUS BUFFER (1102) to signal this condition to the channel.

FIG. 12 illustrates a transmitter circuit which adds the null words as the data is transmitted. The transmission hardware consists of an ASYNCHRONOUS BUFFER (1200) which is similar to the one described in U.S. Pat. No. 5,003,558 of Gregg entitled "Data Synchronizing Buffers for Data Processing Channels". The ASYNCHRONOUS BUFFER (1200) receives data words from the channel over the TRANSMIT DATA lines (1201). Along with the data words, the running exit disparity of the data is sent to the ASYNCHRONOUS BUFFER (1200) over the DISP line (1201). The ASYNCHRONOUS BUFFER (1200) is loaded under control of the SYSTEM CLOCK (1202) from the channel. Data words from the ASYNCHRONOUS BUFFER (1200) are unloaded and sent to the serializer through the MULTIPLEX (1203) over the TRANSMIT DATA OUT lines (1204). The TRANSMIT CLOCK (1205) from the serializer is used to unload the ASYNCHRONOUS BUFFER (1200) at the serializer rate. The TRANSMIT CLOCK (1205) drives the ASYNCHRONOUS BUFFER's UNLOAD line (1206) through an AND circuit (1207). The other input to the AND circuit (1207) comes from the null pacing function described later. When the line from the null pacing function (1208) is on, data words are unloaded from the ASYNCHRONOUS BUFFER (1200) and sent to the MULTIPLEX (1203). The output of the AND circuit (1207) also gates the multiplex (1203) such that data words from the ASYNCHRONOUS BUFFER (1200) are sent on the TRANSMIT DATA OUT lines (1204). When the line from the null pacing function (1208) is off, the UNLOAD (1206) to the ASYNCHRONOUS BUFFER (1200) is suppressed, and the MULTIPLEX (1203) selects a data word from the NULL GEN circuit (1209). The NULL GEN circuit (1209) uses the DISP value (1210) from the previous data word from the ASYNCHRONOUS buffer (1200) to generate the proper NULL WORD (described above).

The null pacing function is shown on the right side of the diagram and consists of a NULL PACE register (1210) which is clocked by the SYSTEM CLOCK (1202) and loaded from the channel over the NULL PACE lines (1211). Data from the NULL PACE register (1210) feeds the MULTIPLEX (1212) which in turn feeds the NULL CNTR (null counter) register (1213). The NULL CNTR register (1213) is clocked by the TRANSMIT CLOCK (1205). The other inputs to the MULTIPLEX (1212) are the outputs of -1 (decrement logic) (1214) and a constant value of '01'x (1215). The MULTIPLEX (1212) is controlled by the LOGIC function (1216). The output of the NULL PACE register (1210) is examined by an 8 input nand circuit shown as =0 (1217). When the NULL PACE register (1210) contains a value of '00'x, the output of =0 (1217) is on. The

output of =0 (1217) is connected to the LOGIC function (1216). Bits 0 through 6 of the NULL CNTR (1213) are connected to a 7 input nand circuit shown as =0 (1218). When the NULL CNTR register (1213) contains a value of '00'x or '01'x, the output of =0 (1218) is on. The output of =0 (1218) is connected to the LOGIC function (1216). Bit 7 of the NULL CNTR register (1213) is also connected to the LOGIC function (1216). The USE NULL PACE=1 line (1212) from the channel feeds the SYNC logic (1220) which is used to synchronize this line to the TRANSMIT CLOCK. The output of the SYNC logic (1220) is connected to the LOGIC function (1216).

The operation of the null pacing function is described by the table in FIG. 13 which shows the equations, in table form, of the LOGIC function (1216). There are three modes of operation.

NORMAL MODE. When the value of NULL PACE (1210) is not '00'x, and the USE NULL PACE=1 line (1219) is off, the LOGIC decrements the value in the NULL CNTR (1213) until it reaches '00'x. When the value in the NULL CNTR (1213) is '00'x, the logic inserts a null word into the data stream by turning off its input to the AND circuit (1207) which gates the TRANSMIT CLOCK (1205). The LOGIC also reinitializes the NULL CNTR (1213) to the value in the NULL PACE register (1210), and the process repeats.

DISABLE NULL PACING. When the value of NULL PACE (1210) is '00'x, and the USE NULL PACE=1 line (1219) is off, the LOGIC constantly gates the TRANSMIT CLOCK (1205) through the AND circuit (1207). No null words are ever inserted into the TRANSMIT DATA OUT lines (1204). USE NULL PACE=1. When this line (1219) is on, the value in the NULL PACE register (1210) is ignored, and null words are inserted into the data stream every other data word. If when the USE NULL PACE=1 (1219) turns on, the value in the NULL CNTR (1213) is greater than '01'x, the LOGIC forces a '01'x value into the NULL CNTR register (1213). This feature starts null pacing of 1 immediately without waiting for a large pacing count to be exhausted. When the value in the NULL CNTR (1213) reaches '00'x, the LOGIC inserts the null word, and reinitializes the NULL CNTR (1213) to '01'x. Null words are thus alternated with data words. Since the contents of the NULL PACE register (1210) are not examined in this mode, the channel is free to change the contents of NULL PACE register (1210) without causing metastability problems.

While we have described our preferred embodiments of our invention, it will be understood that those skilled in the art, both now and in the future, may make various improvements and enhancements which fall within the scope of the claims which follow. These claims should be construed to maintain the proper protection for the invention first disclosed.

What is claimed is:

1. In a data processing system with interconnected elements each of which can function as both a transmitter and a receiver of an information stream, a method for operating the elements at each end of a link wherein each of said elements and the link operates at its own operating rate and wherein said operating rates are independent from one another comprising the steps of:

receiving a frame which contains the maximum rate at which a receiver can receive said information stream;
determining an information rate equivalent to the maximum rate at which both the transmitter and the receiver can process said information stream;
setting the transmitter to transmit non-null words at the determined information rate by inserting null words at

the transmitter, as needed, to reduce the effective rate of non-null words transmitted therefrom in accordance with said determined information rate; and

deleting null words as they are received at the receiver.

2. A method according to claim 1 wherein:

the receiver can specify to the transmitter the maximum rate at which it can receive the information stream.

3. In a data processing system with interconnected elements which can function as both a transmitter and receiver of an information stream, a method for operating the elements at each end of a link at a speed higher or lower than that of the link comprising:

determining a maximum rate at which a transmitter can transmit said information stream, receiving a frame which contains the maximum rate at which a receiver can receive said information stream;

determining an information rate equivalent to the maximum rate at which both the transmitter and receiver can process said information stream;

setting the transmitter to transmit non-null words at the information rate by inserting null words at the transmitter, as needed, to reduce the effective rate of non-null words transmitted in accordance with said determined information rate; and

deleting null words as they are received at the receiver; wherein an overrun indication is set when a received word is overwritten prior to the system reading a previous word from said transmitted information stream.

4. In a data processing system with interconnected elements which can function as both a transmitter and receiver of an information stream, a method for operating the elements at each end of a link at a speed higher or lower than that of the link comprising:

determining a maximum rate at which a transmitter can transmit said information stream, receiving a frame which contains the maximum rate at which a receiver can receive said information stream;

determining an information rate equivalent to the maximum rate at which both the transmitter and the receiver can process said information stream;

setting the transmitter to transmit non-null words at the determined information rate by inserting null words at the transmitter, as needed, to reduce the effective rate of non-null words transmitted in accordance with said determined information rate; and

deleting null words as they are received at the receiver; wherein when inserting said null words in accordance with said determined information rate a null word follows 8B/10B code encoding rules, and said null word for 8B/10B codes comprises the characters k28.5-D21.5-D21.5D16.5, codes (1100000101 1010101010 1010101010 011011010 and 1100000101 1010101010 1010101010 100100101), which are dynamically selected for insertion based on the running disparity of said information stream, whereby transmitters maintain and use the proper version of said null word so that all of the 8B/10B code encoding rules are met.

5. A method according to claim 4 wherein:

the null word has neutral disparity allowing null words to be added and removed without affecting data being transmitted.

6. A method in accordance with claim 4 wherein:

the null word is chosen to be similar to an idle word so that some of the hardware to recognize a null word can also be used to recognize the idle word.

27

7. In a data processing system with interconnected elements which can function as both a transmitter and receiver of an information stream, a method for operating the elements at each end of a link at a speed higher or lower than that of the link comprising:

determining a maximum rate at which a transmitter can transmit said information stream, receiving a frame which contains the maximum rate at which a receiver can receive said information stream;

determining an information rate equivalent to the maximum rate at which both the transmitter and receiver can process said information stream;

setting the transmitter to transmit non-null words at the information rate by inserting null words at the transmitter, as needed, to reduce the effective rate of non-null words transmitted in accordance with said determined information rate; and

28

deleting null words as they are received at the receiver; wherein the receiver can specify to the transmitter the maximum rate at which it can receive the information stream; and

wherein a pacing rate is reset to the slowest allowable rate when a continuous sequence is received;

the receiver's pacing rate is transmitted at the slowest allowable rate;

the transmitter sets up the pacing rate once the receiver's pacing rate is received;

the pacing rate is set at the maximum rate at which both the transmitter and receiver can process said information stream.

* * * * *